

ABSTRACT

An end to end forwarding architecture includes a memory hub having a first ingress interface for receiving packets from a source port. The packets have associated ingress flow identifiers. A second ingress interface outputs the packets to a switch fabric. An ingress controller manages how the packets are queued and output to the switch fabric. The same memory hub can be used for both per flow queuing and per Class of Service (CoS) queuing. A similar structure is used on the egress side of the switch fabric. The end to end forwarding architecture separates per flow traffic scheduling operations performed in a traffic manager from the per flow packet storage operations performed by the memory hub.